



AMENDMENTS TO THE SPECIFICATION


Please replace the paragraph beginning at page 1, line 10, with the following rewritten paragraph:

 This application is related to application serial no. 09/476,846, now pending titled METHOD AND SYSTEM FOR EVENT IMPACT ANALYSIS, attorney docket no. 3882/2, filed January 3, 2000, which application is hereby incorporated herein by reference in its entirety.


Please replace the paragraph beginning at page 4, line 1, with the following rewritten paragraph:

 Furthermore, there is a need for improved methods and systems for helping administrators make decisions about how to prioritize outages and allocate resources in the correction or prevention of service failures. Commonly assigned application serial no. 09/476,846, now pending and U.S. Patent No. 5,872,911, discussed above, describe different systems for determining the impact of a service failure on customers or users. However, these systems do not quantify the impact in a way to provide the administrator with the ability to compare the effects of outages in different, unrelated services in order to prioritize the allocation of resources, or to perform a strict cost/benefit analysis for the allocation of the resources. Improved methods and systems are thus needed to quantify the cost of a service outage in such a way as to allow the cost to be compared to costs of other service outages in services or systems which may differ in type or use.

Please replace the paragraph beginning at page 13, line 12, with the following rewritten paragraph:

 Figs. ~~23-34~~ 23-24 are graphical representations of techniques for predicting the cost of an outage during the outage in accordance with embodiments of the present invention.

Please replace the paragraph beginning at page 15, line 12 to the paragraph beginning at page 16, line 3, with the following rewritten paragraphs:

 The service 20 level or quality data and service usage data are collected through any conventional devices depending upon the particular service 20. For example, referring to Fig. 2, in one embodiment of the invention, a service accessed by a number of end users 22 is monitored by service level monitors 24 as well as usage meter(s) 26. As shown, the usage meter or meters measure the total amount in which the service 20 is used. This measurement will depend upon the nature of the service. For example, in a telecommunication service, the usage meter 26 measures the number of messages being conveyed by the service 20 in an hour or other time period. Such a telecommunication service 20 may include a wireless telephony system using the smart message service 20 protocol, in which the usage meter measures the number of messages conveyed over the system. Similarly, in a computer network, the amount of traffic may be measured over a given period. For a service consisting of an application software program, the usage may be measured in any appropriate manner, such as time spent, number of memory calls over time, number of documents accessed over time, etc.

Please replace the paragraph beginning at page 16, line 3, with the following rewritten paragraph:

OK The service 20 level monitors 24 monitor the status of the service on a polled basis. That is, they poll the service on a regular basis to determine whether the service is working. In some embodiments, this inquiry is binary in nature, resulting in a status of either working or non-working, sometimes referred to herein as a status of good or bad service. Alternatively, the service level monitors 24 may detect several types of service status, and an outage would be defined as desired as a transition from one status in which the service 20 considered to be working properly to a second status in which it is not providing an acceptable level of service 20. Multiple discrete service monitors 24 may be used to monitor the quality of service 20 from different locations. For example, for a network accessible from several remotely located sites, each site would have a service monitor 24 to monitor the quality of service being received at that site, while the traffic on the entire network would be measured by usage meters 26.

Please replace the paragraph beginning at page 16, line 21 through page 17, line 8 with the following rewritten paragraph:

OK One embodiment of the present invention implemented in a computerized network environment is shown in Fig. 3. In this embodiment various activity on a network 30 is monitored by a network monitoring system 32, such as the NETCOOL suite available from Micromuse Inc., as described above. Other commercially available network monitoring systems 32 may be used, as would be understood by one skilled in the art. The network monitoring system 32 works in conjunction with one or more service quality monitors 24, such as Internet Service Monitors, a set of active monitors which work with the NETCOOL suite to provide availability information on Internet service elements including DNS, HTTP, LDAP, POP3,

all cont
RADIUS, SMTP, and TCP Port Monitoring, or NT Service Monitors, which work with the network monitoring system 32 to collect management events from a network operations center regarding Windows NT® environments.

Please replace the paragraph beginning at page 17, line 9 with the following rewritten paragraph:

am
Further, a set of service usage meters 26A, 26B, and 26C work with the network monitoring system 32 to track the amount of activity seen at different locations on the network 30. In particular, in some embodiments service specific usage meters ~~26A~~ 26B track activity on different services 34 available within the network, including, for example, a cache server, an email server, Radius, HTTP/HTML, LDAP, and other applications and services available to end users over the network 30. Other, network edge usage meters ~~26B~~ 26A measure activity in and out of the network 36 including, for example, a usage meter ~~26B~~ 26A detecting activity at a firewall monitoring application which collects security information, a usage meter ~~26B~~ 26A at Cisco Net Flow, a usage meter ~~26B~~ 26A at a PIX firewall, and others as known to those skilled in the art.

Please replace the paragraph beginning at page 18, line 13 with the following rewritten paragraph:

as
The operation of one embodiment of the probable cause engine 16 is now described with reference to the flow chart in Fig. 4 and the graphical diagrams in Fig. 5-13. Referring to Fig. 4, the probable cause engine receives service quality data from service quality monitors and service usage data from the service usage meters, step 60. This received data is

used to detect the occurrence of an outage, step 62. Where no outage is detected, however, the probable cause engine continues to receive service quality data, step 60. When an outage is detected, the probable cause engine defines a time window, sometimes referred to herein as the likely cause window or LCW, using the service quality data, service usage data, or preferably both, step 64. The probable cause engine then retrieves event data from the event object server for events which occurred during the defined likely cause window and for a period prior to the LCW, step 66. The events retrieved are each assumed to have a non-zero probability of having caused the outage, and it is further assumed that at least one of the retrieved events caused the outage. The engine then computes a probability distribution for the events which maps the probabilities that each event caused the detected outage, step 68. These probabilities are then associated with the events and output to the administrator in a suitable form.

Please replace the paragraph beginning at page 28, line 23 to page 29, line 8 with the following rewritten paragraph:

Referring to Fig. 14, as explained above the service quality monitors regularly poll the service for quality level and the service usage meters continuously measure the activity in a service, step 84. If a service outage is detected, step 86, based on the service quality or service usage data using techniques described above, the costing engine predicts the cost of the outage based on the current outage window length, for current outages, and including various extended outage windows for completed outages, step 88. Where no service outage is detected, the service usage meters continue to measure the activity in a service, step 84. The use of window lengths to predict and compute outage costs is described in greater detail below. Once a predicted cost is computed, it is compared to the outage costs determined for other services, step

agcont 90. This comparison may then be used to prioritize outages and allocate resources to the service outages, step 92.

Please replace the paragraph beginning at page ~~28~~²⁹, line ~~23~~²² to page ~~29~~³⁰, line ~~8~~¹⁰ with the following rewritten paragraph:

aid Where no service outage is detected, step 102, the service quality monitors continue to regularly poll the service for quality levels. If a service outage is detected as completed, step 102, using techniques described above for detecting the onset of an outage and its at least partial recovery, an outage window is defined using the service quality data and/or service usage data, step 104. For example, as described above, the outage window may be defined as the period between two service quality readings of good service with one or more bad quality readings in between. The cost of outage engine then defines an extended window which includes the outage window and a second time window, step 106. This extended window is referred to herein as a cost of outage window. The cost of outage window is the time span over which the usage data is considered for any given outage, with the goal being to analyze the usage data over a period corresponding to that which incorporates all the effects of the outage on usage patterns. The cost of outage window thus extends beyond the bounds of the outage window to pick up the recovery in usage, subsequent to the outage.